

B

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

SRI INTERNATIONAL, INC.,)
a California Corporation,)
Plaintiff and)
Counterclaim-Defendant,)
v.) C. A. No.: 04-1199 (SLR)
INTERNET SECURITY SYSTEMS, INC.,)
a Delaware Corporation, INTERNET)
SECURITY SYSTEMS, INC., a Georgia)
Corporation, and SYMANTEC)
CORPORATION, a Delaware Corporation,)
Defendants and)
Counterclaim-Plaintiffs.)

**INVALIDITY EXPERT REPORT OF
L. TODD HEBERLEIN**

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
II.	QUALIFICATIONS	1
III.	METHODOLOGY AND BASES.....	4
IV.	RELEVANT FIELD OF ART AND PERSON OF ORDINARY SKILL IN THE ART.....	5
V.	THE SPECIFICATION OF THE PATENTS-IN-SUIT	5
	A. INVENTORS' DESCRIPTION OF THEIR INVENTION	5
	1. Written description of the patents-in-suit	7
	2. The claims of the '338 patent.....	11
	3. The claims of the '203 patent.....	13
	4. The claims of the '212 patent.....	15
	5. The claims of the '615 patent.....	16
	B. COMPUTER SCIENCE CONCEPTS RELEVANT TO THE PATENTS.....	18
	C. INTRUSION DETECTION METHODOLOGIES.....	21
VI.	CLAIM LANGUAGE.....	22
VII.	LEGAL STANDARDS – ANTICIPATION AND OBVIOUSNESS	23
VIII.	SUMMARY OF OPINIONS	25
IX.	STATE OF THE RELEVANT ART AS OF THE FILING DATE	27
	A. HISTORY OF NETWORK MONITORING AND INTRUSION DETECTION	27
	1. SRI's history in intrusion detection	33
	2. UC Davis's history in intrusion detection.....	36
	3. Government systems' history in intrusion detection	38
	4. Commercial systems' history in intrusion detection.....	39

X. LACK OF NOVELTY.....	39
A. PRIOR ART PUBLICATIONS AND SYSTEMS	40
1. NSM.....	40
2. DIDS	49
3. ISM	57
4. GrIDS	61
5. JiNao	69
6. EMERALD 1997	76
a. CMAD and Conceptual Overview.....	84
b. Live Traffic Analysis and Conceptual Design 1997.....	85
c. Emerald 1997, Intrusive Activity 1991, NIDES 1994	85
7. Network NIDES.....	86
8. Frank Feather Thesis.....	90
9. HP Open View and RFCs	96
10. NetStalker	104
11. NetRanger	106
12. ISS RealSecure.....	108
13. Network Level Intrusion Detection	109
14. '750 Thompson patent	110
15. Stake Out.....	111
16. Automated Alarm System.....	112
XI. OBVIOUSNESS	113
A. SECONDARY CONSIDERATIONS	119
XII. INVENTORS' FAILURE TO DISCLOSE BEST MODE.....	121
A. LEGAL STANDARD.....	121

B.	DOCUMENTS REL UPON AND METHODOLOGY	121
C.	THE EXPERT SIGNATURE ENGINE, SRI'S BEST MODE FOR ANALYZING NETWORK TRAFFIC DATA, WAS NOT DISCLOSED	123
1.	Expert systems	124
2.	SRI's P-BEST expert system.....	124
3.	SRI had developed a fully-functional version of eXpert prior to November 9, 1998.....	125
a.	P-BEST inference engine and translator.....	127
b.	Integration between P-BEST and other parts of eXpert	127
c.	Suite of P-BEST signatures for analyzing network traffic data.....	128
4.	eXpert was SRI's best mode for practicing signature detection upon network traffic data	130
5.	eXpert was SRI's best mode for analyzing network traffic data	130
6.	In the alternative, the combination of eXpert and estat was the inventors' best mode for analyzing network traffic data	131
7.	eXpert was not disclosed in the patents-in-suit.....	132
D.	ETCPGEN, SRI'S BEST MODE FOR PERFORMING NETWORK MONITORING, WAS NOT DISCLOSED.....	134
E.	eRESOLVE, SRI'S BEST MODE FOR PERFORMING RESPONSE, WAS NOT DISCLOSED	135
XIII.	ENABLEMENT AND SUFFICIENT WRITTEN DESCRIPTION OF THE PATENTS-IN-SUIT.....	136
A.	LEGAL STANDARD.....	136
B.	ANALYSIS OF ENABLEMENT / WRITTEN DESCRIPTION	136
XIV.	PUBLIC AVAILABILITY OF CERTAIN DOCUMENTS.....	139

I. INTRODUCTION

1. I, L. Todd Heberlein, am the President of Net Squared, Inc.
2. I have been retained by counsel for Symantec Corporation as an expert witness in this action. If called to testify as to the truth of the matters stated herein, I could and would do so competently.

II. QUALIFICATIONS

3. I received a Bachelor of Science degree in Computer Science and Math from the University of California, Davis in 1988, and a Masters of Science degree in Computer Science from the University of California, Davis in 1991.
4. I was the primary developer of the first network-based intrusion detection system, the Network Security Monitor (NSM), at UC Davis in the late 1980s and early 1990s. The NSM processed network packets and applied anomaly and signature detection techniques to detect intrusive activity. The work began in 1988, and we published numerous papers describing the work in 1990, 1991, and 1994. By the mid-1990s the NSM was deployed at numerous organizations including UC Davis, the Air Force, the Department of Energy, NASA, and the Department of Justice. The Air Force deployed it at over 100 Air Force sites globally. Within the Air Force and at Lawrence Livermore National Laboratory, the NSM was usually deployed at the organization's gateway to the rest of the Internet.

5. I was also one of the primary developers for the first hierarchical and distributed intrusion detection system, the Distributed Intrusion Detection System, (DIDS) which integrated NSM, host monitors for SunOS and VMS, and a centralized Director. DIDS was deployed at UC Davis and Lawrence Livermore National Laboratory in 1992. I have been told that the Air Force deployed DIDS at other locations at a later date.

6. From 1993 to 1995 I worked on a DARPA contract called Intrusion Detection for Large Networks, and in 1995 I presented results, including an initial fully distributed intrusion detection capability, to our DARPA Program Manager, Teresa Lunt. That work eventually became the Graph-based Intrusion Detection System (GrIDS). In addition to intrusion detection technology, I researched issues about the fundamental nature of computer vulnerabilities, and that work won a “best paper” award at the National Information Systems Security Conference in 1996.¹ The paper was also republished in a book by Dorothy and Peter Denning.²

7. I founded Net Squared, Inc. in 1996. At Net Squared I performed research and development in computer security for numerous organizations including the Air Force, Lawrence Livermore National Laboratory, the Defense Advanced Research Projects Agency (DARPA), Office of Naval Research, the Intelligence Community, and the Federal Aviation Administration. At Net Squared I developed the Network Monitoring Framework (NMF), a library of network monitoring C++ objects, and Network Radar, a suite of network monitoring applications built on the NMF libraries. Network Radar tools were integrated into larger, hierarchical intrusion detection systems by the Air Force and Boeing, including EPIC, EPIC2, AIDE, AFED, and IDIP. Network Radar was deployed at UC Davis, the Air Force’s Rome Research Laboratory, and during numerous Air Force exercises. From approximately 2002-2003 I also led a project called TrendCenter, which integrated and correlated intrusion detection alerts from unrelated organizations. As part of that effort I developed SANS’ initial Internet Storm Center prototype.

¹ L. T. Heberlein, M. Bishop, “Attack Class: Address Spoofing,” 19th National Information Systems Security Conference, Baltimore, MD, 22-25 Oct. 1996, pp. 371-377.

² D. Denning and P. Denning, INTERNET BESIEGED, COUNTERING CYBERSPACE SCOFFLAWS, 1st ed. (Oct. 3, 1997) at Chap. 10.

8. In addition to traditional intrusion detection capabilities (anomaly detection and signature detection), I developed or help develop several other intrusion detection technologies. For example, we profiled network services in NSM and DIDS, and that work was also part of a feature selection effort led by Jeremy Frank.³ The profiling of network services work, later based on a feed-forward back-propagation neural network, was also rolled into Network Radar's Non-Cooperative Service Recognition (NCSR) technology.

9. I also did the initial work on network thumbprinting that was introduced in the 1992 Internetwork Security Monitor paper,⁴ and I helped refine the thumbprinting technology that was published in a 1995 IEEE paper.⁵ Thumbprinting uses a multivariate statistical technique called principal component analysis to reduce a high dimensional object to a small dimensional object. Vector distances of the lower dimensional objects were then used to determine if the objects were correlated.

10. Both the NCSR and the thumbprinting work were also rolled into Network Radar, which was able to detect the earliest stages of the ILOVEYOU worm as it hit the Air Force's Rome Labs.

11. A summary of my professional experience and publications are attached as Exhibit A.

12. I receive compensation in the amount of \$258.00 per hour for the time that I devote to this matter. My compensation is not dependent in any way on the outcome of this matter.

³ J. Frank , "Machine Learning and Intrusion Detection: Current and Future Directions," Proc. of the 17th National Computer Security Conference, October 1994.

⁴ L.T. Heberlein, B. Mukherjee, K.N. Levitt, "Internetwork Security Monitor: An Intrusion-Detection System for Large-Scale Networks," Proc. 15th National Computer Security Conference, pp. 262-271, Oct. 1992.

⁵ S. Staniford-Chen, and L.T. Heberlein, "Holding Intruders Accountable on the Internet," Proc. of the 1995 IEEE Symposium on Security and Privacy, Oakland, CA, 8-10 May 1995, pp. 39-49.

III. METHODOLOGY AND BASES

13. In preparing my opinions and analysis of the art, I have thoroughly reviewed the entire specification and claims of U.S. Patents No. 6,321,338 (the '338 patent); 6,484,203 (the '203 patent); 6,708,212 (the '212 patent); and 6,711,615 (the '615 patent) (collectively, the patents-in-suit). I have also reviewed each of the prosecution histories and the Microfiche Appendix included with the patents-in-suit.

14. I have also reviewed the expert reports of Mr. Daniel Teal and Mr. Frederick Avolio, and have indicated in my report instances where I have relied upon these reports.

15. I have reviewed an extensive body of prior art publications and product documentation. I have also spoken directly with a number of individuals who I understand and believe to be personally knowledgeable with respect to the prior art embodiments discussed below. A list of the prior art publications and documentation I have reviewed and the individuals with whom I have spoken in forming the opinions set forth below is attached as Exhibit B.

16. I have also compared each of the claims of the patents-in-suit with certain prior art publications and embodiments discussed below, and charts comparing the claims of the patents with those printed publications and embodiments are attached as Exhibits C-V.

17. In general, my methodology of inquiry with respect to each of the principal prior art publications and embodiments relied upon in my opinions was as follows:

- a. With regard to prior art publications, I reviewed each publication in detail.
- b. With regard to product embodiments, I reviewed the marketing literature, product documentation and manuals relating to the prior art system to familiarize myself with the features, functions and capabilities of the system.
- c. With regard to product embodiments, I typically also engaged in a lengthy conversation with at least one individual who was personally

knowledgeable of the facts relating to the development, marketing and history of the prior art embodiment. In this conversation I attempted both to challenge and to confirm my understanding of the facts that I had derived from my prior investigations.

- d. In general I have attempted to maintain a skeptical perspective with respect to the materials I have reviewed in order to satisfy myself that the bases for my opinions are well-founded and truly convincing to one skilled in the art.

IV. RELEVANT FIELD OF ART AND PERSON OF ORDINARY SKILL IN THE ART

18. The patents-in-suit relate to the field of network monitoring, in particular network monitoring for the purpose of detecting suspicious or intrusive network activity. I have based this determination upon my examination of the patents-in-suit. For instance, the '338, '212 and '615 patents are all entitled "Network Surveillance." The '203 patent is entitled "Hierarchical Event Monitoring and Analysis." The '338 patent claims a "method of network surveillance." The '203, '212 and '615 patents claim an "enterprise network monitoring system." The patents-in-suit claim network monitoring for the purpose of detecting "suspicious network activity."

19. In light of my opinion as to the field of art relevant to the patents-in-suit, I am of the opinion that one of ordinary skill in the art as of the filing date would have been someone with an undergraduate degree in Computer Science with at least three to five years experience in computer programming and network design with an emphasis in network monitoring technology and intrusion detection.

V. THE SPECIFICATION OF THE PATENTS-IN-SUIT

A. Inventors' description of their invention

20. I have reviewed the specifications of all four of the patents-in-suit, including the written description and the claims. I have also reviewed the file histories of all four patents-in-suit. The patents-in-suit also include a Microfiche Appendix, which I

reviewed in detail. I have based the following description of the alleged inventions on these documents.

21. The '203, '212 and '615 patents are all continuations of the '338 patent. I understand that this means that all four patents-in-suit are entitled to the filing date of the '338 patent, which is November 9, 1998. I also understand that as continuations, the '203, '212 and '615 patents may not add "new matter" or additional disclosures to the specification of the '338 patent. Although the "References Cited," "Abstract" and "Summary" sections of the four patents-in-suit differ in some respects, the Figures and "Detailed Description" sections are essentially identical and provide a common description of the alleged inventions. Because the most relevant portions of the written description of the four patents-in-suit share a common disclosure, I have cited only to the '338 patent in my description of the alleged inventions. However, this description applies to all four of the patents-in-suit.

22. In addition, I understand that the patents' specification purports to incorporate-by-reference the information in A. Valdes and D. Anderson, "Statistical Methods for Computer Usage Anomaly Detection Using NIDES," Proceedings of the Third International Workshop on Rough Sets and Soft Computing, January 1995 ("Statistical Methods").⁶ However, I also understand that SRI has admitted⁷ that this publication is not "essential material," which is defined in the Manual of Patent Examining Procedures Section 608.01 (7th Ed., July 1998) as:

"Essential material" is defined as that which is necessary to (1) describe the claimed invention, (2) provide an enabling disclosure of the claimed invention, or (3) describe the best mode (35 U.S.C. 112).

⁶ '338 col. 5:43-49.

⁷ See SRI International, Inc.'s Responses to Defendant Symantec Corporation's First Set of Requests for Admission [RFA No. 2].

23. Because this publication cannot be relied upon to describe or enable the alleged inventions, I have not included a discussion of it in this description.

24. Furthermore, I also understand that the patents' specification purports to incorporate-by-reference the information in P. Porras and A. Valdes, "Live Traffic Analysis of TCP/IP Gateways," Networks and Distributed Systems Security Symposium, March 1998 ("Live Traffic Analysis").⁸ However, I also understand that SRI has admitted that this publication is not "essential material" as well.⁹ Because this publication cannot be relied upon to describe or enable the alleged inventions, I have not included a discussion of it in this description.

1. Written description of the patents-in-suit

25. The common written description of the patents-in-suit describes a hierarchical scheme for the monitoring and analysis of networks for the purpose of intrusion detection. An enterprise network is shown that contains different "domains."¹⁰ Each domain "includes one or more computers offering local and network services that provide an interface for requests internal and external to the domain."¹¹

26. The enterprise network includes a set of "network monitors" for analyzing network activity.¹² These monitors are deployed in an "analysis hierarchy"¹³ and include lowest-level "service monitors," as well as "domain monitors" and "enterprise monitors." The service monitors analyze network traffic / network packets handled by "network entities" such as gateways, routers, or firewalls.¹⁴ The service monitors produce reports

⁸ '338 col. 12:61-65.

⁹ See SRI International, Inc.'s Responses to Defendant Symantec Corporation's First Set of Requests for Admission [RFA No. 4].

¹⁰ See '338, Fig. 1 and accompanying description.

¹¹ '338 col. 3:17-20.

¹² '338 col. 3:32-35.

¹³ '338 col. 3:33-34.

¹⁴ '338 col. 3:42-45.

and disseminate them to other monitors via a subscription-based distribution scheme.¹⁵ Domain monitors receive reports from service monitors and distribute their own reports to enterprise monitors.¹⁶ In turn, enterprise monitors correlate reports across their set of monitored domains.¹⁷

27. The written description states that both domain monitors and enterprise monitors “correlate” reports. However, no description of how correlation is performed is provided. The written description does not contain any algorithms for correlating reports, nor does it provide detail on what types of reports may be correlated. The written description does suggest that enterprise monitors should focus upon recognizing commonalities in reports such as “the spreading of a worm or a mail system attack repeated throughout the enterprise” but does not identify specific events or reports that may identify such scenarios.¹⁸

28. Figure 2 of the patents-in-suit illustrates a monitor. All of the monitors (service, domain and enterprise) use the same monitor code-base.¹⁹ Monitors are configured for different tasks via a “resource object” shown in Figure 3 of the patents-in-suit. The resource object is used to “tun[e] the generic monitor code-base to a specific event stream.”²⁰

29. Each monitor can analyze “event records that form an event stream.”²¹ Event records are created from network traffic / network packets in a variety of different ways, and packets may be selected for analysis based upon different criteria.²² Each monitor includes one or more analysis engines. Event records are forwarded to the

¹⁵ ‘338 col. 3:55-65, *see also* col. 9:22-11:3.

¹⁶ ‘338 col. 3:66-4:18.

¹⁷ ‘338 col. 4:18-47.

¹⁸ ‘338 col. 4:38-39.

¹⁹ ‘338 col. 11:4-11.

²⁰ ‘338 col. 11:27-28.

²¹ ‘338 col. 4:61-62.

²² ‘338 col. 4:61-5:5.

analysis engines within a monitor for analysis.²³ “The analysis engines 22, 24 receive large volumes of events and produce smaller volumes of intrusion or suspicion reports that are then fed to the resolver.”²⁴

30. Two different types of analysis engines are described: a “signature analysis engine” and a “statistical profiling,” which perform different types of analysis on the data collected by the monitors.²⁵ The signature engine can perform threshold analysis, which detects when the number of occurrences of a specific event exceeds a preset level.²⁶ The signature engine uses abstract representations of known attacks to compare against incoming event streams.²⁷

31. The statistical engine performs statistical profile-based anomaly detection. In contrast to the signature engine, the algorithms used for the statistical profiling analysis require no prior knowledge of intrusive or exceptional activity.²⁸ The statistical engine uses “statistical measures to profile network activity indicated by an event stream.”²⁹ Statistical measures are variables created from event records.³⁰ Four types of measures are described: categorical, continuous, intensity, and event distribution.³¹

²³ ‘338 col. 5:34-35.

²⁴ ‘338 col. 8:13-15.

²⁵ See ‘338 Fig. 2 and accompanying description at 4:48-53. Although the patents state “a monitor 16 may include additional analysis engines that may implement other forms of analysis” I have not found a description of any other type of analysis engine in the patents’ written description.

²⁶ ‘338 col. 7:24-26, 7:45-55.

²⁷ ‘338 col. 7:23-42.

²⁸ ‘338 col. 6:57-58.

²⁹ ‘338 col. 5:36-38.

³⁰ “The profile engine 22 can use a wide range of multivariate statistical measures to profile network activity indicated by an event stream.” ‘338 col. 5:36-38. “The monitor builds a statistical model of network activity from the network packets, for example, by building 68 long-term and short-term statistical profiles from measures derived from the network packets.” ‘338 col. 12:47-53.

³¹ ‘338 col. 5:49-6:37.

These measures are used to create a long-term statistical profile and a short-term statistical profile.³²

32. While the long-term statistical profile characterizes historical activity, the short-term statistical profile “characterizes recent activity.”³³ The patents explain that both the long-term and the short-term profiles are “aged to adapt to changes in subject activity.”³⁴ The short-term profile is periodically used to update the long-term profile and is then cleared. The short-term profile is compared to the long term profile to determine if recent activity is anomalous.³⁵

33. A monitor also includes a resolver for implementing a response policy (the patents also use the term “countermeasure” to mean the same thing as response).³⁶ The resolver invokes either passive or active types of responses, or countermeasures, such as report dissemination to other monitors or severing a communications channel.³⁷ A monitor also includes an application programmers’ interface (API) to “encapsulat[e] monitor functions” and “eas[e] integration of third-party intrusion-detection tools.”³⁸

34. The appendix to the patents-in-suit contains portions of the source code for the estat program. Estat was the inventors’ embodiment for statistical profiling.³⁹ The estat code in the appendix is “generic” and not specific to reading network traffic – instead, the code would operate on any data provided in the correct format. The appendix also contains code for some basic infrastructure plumbing for communications between various code modules. In my opinion, the estat code contains no further description of

³² ‘338 col. 6:38-50.

³³ ‘338 col. 6:44-47.

³⁴ ‘338 col. 6:38-52.

³⁵ ‘338 col. 6:38-7:3.

³⁶ ‘338 col. 4:55-56; 11:57 – 12:19.

³⁷ ‘338 col. 8:19-21; col. 12:7-19.

³⁸ ‘338 col. 4:57-60.

³⁹ Porras 30(b)(6) Tr. 171-173.

how to implement the alleged inventions other than to illustrate minor unclaimed software implementation details not relevant to the claimed inventions.⁴⁰

2. The claims of the '338 patent

35. I understand that it is the claims of the patents-in-suit that define the alleged inventions. In the following sections I have attempted to distill each patent's claims down to a representative claim for each patent, and a discussion of the additional limitations in other claims of each patent.

36. I understand that SRI has asserted '338 claims 1-8, 10-19, 21, 24 and 25 against Symantec in this litigation.⁴¹ Claims 1, 21, 24 and 25 are all independent claims.

37. Independent claims 1, 21 and 25 of the '338 patent are directed to a "method of network surveillance" and independent claim 24 is directed to a "computer program product." Although claim 1 is a method claim and claim 24 is an apparatus claim, these two claims contain virtually identical limitations. Thus, claim 1 is representative of the alleged invention claimed:

1. A method of network surveillance, comprising:
receiving network packets handled by a network entity;
building at least one long-term and at least one short-term statistical profile from at least one measure of the network packets, the at least one measure monitoring data transfers, errors, or network connections;
comparing at least one long-term and at least one short-term statistical profile; and

⁴⁰ For example, one of the inventors, Mr. Valdes, testified that the patent code illustrates "software engineering aspects" such as speed and performance that would not be disclosed in the patents' specification. Valdes Tr. 349. The patents-in-suit do not require any particular speed or performance by the claimed system or method.

⁴¹ Although SRI has not asserted the remaining '338 patent claims against Symantec, these claims are substantially similar to other asserted claims and I believe that these claims are similarly invalid for the reasons discussed in my report.

determining whether the difference between the short-term statistical profile and the long-term statistical profile indicates suspicious network activity.

38. Independent claim 21 claims building “multiple short-term statistical profiles” but otherwise does not differ substantively from claim 1. Independent claim 25 requires “receiving packets at a virtual private network entity” and instead of separate “comparing” and “determining” limitations, claims “comparing at least one long-term statistical profile with at least one short-term statistical profile to determine whether the packets indicate suspicious network activity.”

39. Dependent claims 2-4 require particular types of “data transfers” be monitored:

- network packet data transfer commands (2),
- network packet data transfer errors (3),
- network packet data transfer volume (4).

40. Dependent claims 5-7 require particular types of “network connections” be monitored:

- network connection requests (5),
- network connection denials (6),
- a correlation of network connections requests and network connection denials (7).

41. Dependent claim 8 requires a particular type of “error” be monitored (error codes included in a network packet), and dependent claim 10 further requires that the “error code” of claim 8 comprise “an error code indicating a reason a packet was rejected.”

42. Dependent claims 11-12 and 16-17 require various different types of response:

- responding (11),
- responding comprises transmitting an event record to a network monitor (12),
- responding comprises altering analysis of the network packets (16),
- responding comprises severing a communication channel (17).

43. Dependent claim 13 requires “transmitting the event record to a hierarchically higher network monitor.”

44. Dependent claim 14 requires “transmitting the event record to a network monitor that receives event records from multiple network monitors.”

45. Dependent claim 15 requires that the network monitor of claim 14 “correlates activity in the multiple network monitors based on the received event records.”

46. Dependent claim 18 requires that the network packets comprise TCP/IP packets.

47. Dependent claim 19 requires that the network entity comprise a gateway, a router, or a proxy server.

3. The claims of the ‘203 patent

48. I understand that SRI has asserted ‘203 claims 1-22 (all of the claims) against Symantec in this litigation.

49. The ‘203 patent contains two independent claims: claim 1 to a “method of hierarchical event monitoring and analysis within an enterprise network” and claim 12 to an “enterprise network monitoring system.”⁴² Although claim 1 is a method claim and

⁴² See ‘203 patent claims 1 and 12.

claim 12 is a system claim, both contain virtually identical limitations. Thus, claim 1 is representative of the alleged invention claimed:

1. A computer-automated method of hierarchical event monitoring and analysis within an enterprise network comprising:

deploying a plurality of network monitors in the enterprise network;

detecting, by the network monitors, suspicious network activity based on analysis of network traffic data selected from the following categories: {network packet data transfer commands, network packet data transfer errors, network packet data volume, network connection requests, network connection denials, error codes included in a network packet};⁴³

generating, by the monitors, reports of said suspicious activity; and

automatically receiving and integrating the reports of suspicious activity, by one or more hierarchical monitors.

50. Both claim 1 and claim 12 each have ten similar dependent claims, requiring:

- “correlating intrusion reports reflecting underlying commonalities” (2 and 13),
- “invoking countermeasures” (3 and 14),
- “network monitors include an API for encapsulation of monitor functions and integration of third-party tools (4 and 15),
- “the enterprise network is a TCP/IP network” (5 and 16),
- “network monitors are deployed at one or more of... {gateways, routers, proxy servers}” (6 and 17),
- deploying a “plurality of service monitors” among “multiple domains” (7 and 18),
- “receiving and integrating is performed by a domain monitor” (8 and 19),

⁴³ I understand that in order to invalidate this limitation, a prior art reference need only disclose one of the claimed “network traffic data” categories.

- “deploying a plurality of domain monitors within the enterprise network” (9 and 20),
- “receiving and integrating is performed by an enterprise monitor” (10 and 21), and
- wherein the domain monitors “establish peer-to-peer relationships with one another” (11 and 22).

4. The claims of the ‘212 patent

51. I understand that SRI has asserted ‘212 claims 1-24 (all of the claims) against Symantec in this litigation.

52. The ‘212 patent contains two independent claims: claim 1 to a “method for monitoring an enterprise network” and claim 14 to an “enterprise network monitoring system.”⁴⁴ Although claim 1 is a method claim and claim 14 is a system claim, both contain virtually identical limitations. Thus, claim 1 is representative of the alleged invention claimed:

1. Method for monitoring an enterprise network, said method comprising the steps of:
 deploying a plurality of network monitors in the enterprise network;
 detecting, by the network monitors, suspicious network activity based on analysis of network traffic data, wherein at least one of the network monitors utilizes a statistical detection method;
 generating, by the monitors, reports of said suspicious activity; and
 automatically receiving and integrating the reports of suspicious activity, by one or more hierarchical monitors.

53. Claims 1 and 14 of the ‘212 patent are very similar to the ‘203 independent claims, except that claims 1 and 14 of the ‘212 patent require only “analysis of network traffic data” without explicitly listing categories of such network traffic data.

⁴⁴ See ‘212 patent claims 1 and 14.

In addition, the '212 claims 1 and 14 require that "at least one of the network monitors utilizes a statistical detection method."

54. Both claim 1 and claim 14 each have ten similar dependent claims, whose limitations are identical to the dependent claim limitations listed above for the '203 dependent claims. In addition, '212 independent claim 1 contains two additional dependent claim limitations:

- "at least one of the network monitors utilizes a signature matching detection method" (2), and
- "wherein the monitor utilizing a signature matching detection method also utilizes a statistical detection method" (3).

5. The claims of the '615 patent

55. I understand that SRI has asserted '615 claims 1-23, 34-53, 64-73, and 84-93 against Symantec in this litigation.⁴⁵ Claims 1, 13, 34, 44, 64, and 84 are all independent claims.

56. Claim 1 to a "method of hierarchical event monitoring and analysis within an enterprise network" and claim 13 to an "enterprise network monitoring system"⁴⁶ both contain virtually identical limitations. Independent claims 34, 44, 64 and 84 all contain similar limitations to independent claims 1 and 13, with minor variations discussed below. Thus, claim 1 is representative of the alleged invention claimed:

1. A computer-automated method of hierarchical event monitoring and analysis within an enterprise network comprising:
deploying a plurality of network monitors in the enterprise network;

⁴⁵ Although SRI has not asserted the remaining '615 patent claims against Symantec, these claims are substantially similar to other asserted claims and I believe that these claims are similarly invalid for the reasons discussed in my report.

⁴⁶ See '615 patent claims 1 and 13.

detecting, by the network monitors, suspicious network activity based on analysis of network traffic data selected from one or more of the following categories: {network packet data transfer commands, network packet data transfer errors, network packet data volume, network connection requests, network connection denials, error codes included in a network packet, network connection acknowledgements, and network packets indicative of well-known network-service protocols};⁴⁷

generating, by the monitors, reports of said suspicious activity; and

automatically receiving and integrating the reports of suspicious activity, by one or more hierarchical monitors.

57. Claims 1 and 13 of the '615 patent are virtually identical to claims 1 and 12 of the '203 patent, except that the '615 patent claims include two additional "network traffic data" categories: "network connection acknowledgments" and "network packets indicative of well-known network-service protocols."

58. Unlike claims 1 and 13 of the '615 patent, independent claims 34, 44, 64 and 84 require only "analysis of network traffic data" without listing particular categories of network traffic data. Each of these independent claims also includes an additional limitation:

- "at least one of the network monitors is deployed at a gateway" (34),
- "at least one of the network monitors is deployed at a router" (44),
- "at least one of the network monitors is deployed at a firewall" (64), and
- "at least one of the network monitors is deployed at one or more of the following facilities of the enterprise network: {gateways, routers, proxy servers, firewalls}" (84).

59. Claim 1 has eleven dependent claims (2-12), claim 13 has ten dependent claims (14-23), and claims 34, 44, 64 and 84 each have nine dependent claims. None of

⁴⁷ I understand that in order to invalidate this limitation, a prior art reference need only disclose one of the claimed "network traffic data" categories.

these dependent claims contain any new limitations not already encompassed by the '203 and '212 claims.

B. Computer science concepts relevant to the patents

60. The patents-in-suit are targeted towards one of skill in the art of computer science, with some understanding of mathematical and statistical algorithms used in intrusion detection. Furthermore, the prior art discussed in this report is similarly targeted towards one of skill in the art of computer science. In order to be able to present the context of this work to a jury and explain the functionality of these systems, I may present at trial an explanation of certain important concepts in computer science.

61. These topics may include: the fundamentals of important networking protocols such as IP, TCP, UDP, FTP, HTTP, Telnet, and SNMP; what a network packet or datagram is; how a packet sniffer works; how TCP connections work; and the benefits of using network packets as a data stream (including leveraging widespread use of Internet Standards or RFCs). I may also offer testimony on general network architectures including distributed architectures, peer-to-peer communications, client-server communications, and communications with remote agents. I may also offer testimony about the basic features and functionality of certain network infrastructure such as clients, servers, routers, firewalls, packet filters, proxy servers, gateways, bridges, switches, virtual private networks (VPNs) and end-nodes.

62. Networked computers typically communicate using several standard network protocols. A protocol is a set of standard rules for data representation, signaling, data exchange and error detection. These protocols are typically depicted as being organized as a stack.⁴⁸ One of the key benefits of this stack-based design is that a

⁴⁸ See W. Richard Stevens, "TCP/IP Illustrated, Volume 1—The Protocols" (Addison-Wesley 1994) (TCP/IP Illustrated, Volume 1) at 2 (depicting a four-layer stack). Another

protocol at a given layer need not be concerned with protocols at layers not directly adjacent to it. For example, TCP, which is designed to be used with IP, need not be concerned with which link layer protocol is being used to transport IP data. This decoupling of protocols enables communications between heterogeneous networks. I am familiar with several protocols in each layer of the network, and at trial I may testify as to the operation of these protocols.

63. Network communications typically occur via network packets. A network packet is a discrete unit of data crossing a network. Several network protocols—including TCP, IP, UDP, and ICMP—specify how to format network packets.⁴⁹ Network packets typically include a header portion and a data portion. An important concept in communicating via network packets is “encapsulation.”⁵⁰ Typically, a packet of information formatted according to a protocol at a given layer is encapsulated within a packet formatted according to a protocol in the layer directly beneath it. For example, in order for an FTP client to send a command to an FTP server, the FTP command is encapsulated within the data portion of the TCP packet. The TCP packet is then encapsulated in the data portion of an IP packet. The IP packet is then encapsulated in the data portion of a packet formatted according to a link-layer protocol such as Ethernet. Thus, as information—such as an FTP command—is passed down the stack of protocols, each protocol in turn encapsulates the information it receives. It is this encapsulation that enables the decoupling of protocols mentioned earlier.

64. The opposite operation—demultiplexing—occurs when the encapsulated network packet is received.⁵¹ When the IP packet is received by the server, the IP

common depiction is the “OSI 7 Layer Model.” See “OSI Seven-Layer Model,” available at <http://www.freesoft.org/CIE/Topics/15.htm>.

⁴⁹ See TCP/IP Illustrated Volume 1 at 34 (IP packet format), 70 (ICMP packet format), 145 (UDP packet format), 255 (TCP packet format),

⁵⁰ See TCP/IP Illustrated Volume 1 at 9-10.

⁵¹ See TCP/IP Illustrated Volume 1 at 11.

implementation on the server will extract the TCP packet from the IP packet by extracting the data portion of the IP packet and then send this extracted TCP packet to the TCP implementation. The TCP implementation will then likewise extract the FTP command from the TCP packet and send it to the FTP implementation.

65. Because of the stack-based relationship between the protocols, concepts at one layer often have analogs at other layers. For example, the concept of an error has meaning at several layers of the stack. In the context of communications between an FTP client and an FTP server, the FTP protocol specifies error messages that are sent in response to syntax errors in FTP commands.⁵² But FTP communications can also generate errors at other layers. If an FTP client sends a FTP command to an invalid IP address, an ICMP packet containing a “host unreachable” error code will be generated and sent to the client. Further down the stack, the transmission of an FTP command may cause a packet collision at the link layer, which is akin to an error. Likewise the concept of a session has meaning at multiple layers of the network stack. A session is a related communication between two systems for a given period of time.

66. By 1997, the use of standard network protocols was already widespread. The prevalence of the TCP/IP protocol suite enabled communications between heterogeneous network devices. In other words, the TCP/IP protocol suite and related standard protocols served as the lingua franca of computer networks, such that a device wanting to communicate with heterogeneous devices on the network need only implement the standard network protocols.

67. The nature of the TCP/IP protocol suite made network packets an ideal subject of analysis for intrusion detection systems.⁵³ By analyzing network packets, an

⁵² See TCP/IP Illustrated Volume 1 at 424.

⁵³ See Mukherjee, Heberlein, and Levitt, “Network Intrusion Detection,” IEEE Network May/June 1994 26-41, 33.

intrusion detection system could analyze a wide range of heterogeneous network entities without needing to understand the various entity-specific audit data formats. In other words, the analysis of network packets enabled intrusion detection systems to scale to monitor larger and more heterogeneous groups of computers.

68. Packets travel over the network via a series of hops from one networked computer to another. Thus, packets sent between two remote computers typically pass through several intermediate network entities. Some of these entities specialize in moving and routing packets across the network. A router is an example of such a specialized entity.

69. By 1997, the vast body of known malicious network-based attacks included attacks that targeted many of the protocol layers as well as several different types of network entities. At trial, I may testify regarding the operation of these attacks and efforts to thwart them.

C. Intrusion detection methodologies

70. There are a variety of different methods for detecting computer intrusions. For example, the specification of the patents-in-suit refers to statistical profiles, statistical anomaly detection, signature matching detection, and threshold analysis. In addition, the intrusion detection field also uses terms such as "misuse detection," "rule-based detection," "expert systems," "heuristic detection," "specification based detection," "protocol anomaly detection," and numerous other terms. Often, the definitions for such terms are not well-defined in the field, or may occasionally be used to mean slightly different things.

71. In general, I tend to characterize methods for detecting computer intrusions into two broad classes: anomaly-detection systems and misuse detection systems. One distinction that can be made between these two classes is that "anomaly

“detection” typically does not require prior knowledge of what an attack may look like, but instead watches for deviations from normal behavior and flags it as anomalous, without necessarily having knowledge of whether the behavior is actually “bad” or not. Depending on the sensitivity of the system, this can lead to false positives on behavior that deviates from normal, but is not actually malicious. However, anomaly detection may in theory be able to detect “new” attacks that do not match any known attack pattern. By contrast, misuse-detection systems typically rely on some form of known bad or good behavior that they can try to detect.

72. However, there are many different ways to characterize and categorize the different analysis methods that may be used for intrusion detection. At trial I may testify about some or all of these different analysis methods.

VI. CLAIM LANGUAGE

73. I understand that the Court has not yet construed certain claim terms of the claims of the patents-in-suit. Since the Court has not yet issued a decision construing the claims of the patents-in-suit, I have been asked, for purposes of this analysis, to assume that the Court adopts the claim construction positions advanced by SRI. In certain instances I have also included a supplemental analysis in the event the Court adopts the claim construction positions advanced by Symantec or ISS.

74. I understand that claim language is generally construed in accordance with its ordinary and customary meaning to those skilled in the relevant art as of the date of the invention. I also understand that claim terms should be given the meaning that is objectively discerned from the specification and/or the prosecution history by one of ordinary skill in the art as of the date of the invention, even if that meaning differs from the term’s ordinary and customary meaning. Where no party has offered a construction for a particular claim term, I have relied upon the ordinary meaning of that term.

VII. LEGAL STANDARDS – ANTICIPATION AND OBVIOUSNESS

75. I understand that a party challenging the validity of an issued United States patent bears the burden of proving invalidity by clear and convincing evidence.

76. I understand that a patent claim is presumed valid unless it is shown by clear and convincing evidence that the claimed subject matter was not novel as of the date of invention, or even if novel, that the differences between the subject matter claimed and the prior art were such that the subject matter as a whole would have been obvious to a person having ordinary skill in the art to which the subject matter pertains.

77. I understand that a claimed invention is not novel if:

- a. the invention was known or used by others in this country, or patented or described in a printed publication in this or a foreign country, before the invention thereof by the applicant for patent, or
- b. the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of the application for patent in the United States, or
- c. the invention was described in a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent.⁵⁴

78. I understand that lack of novelty requires an invalidating prior art reference to disclose each and every limitation of the claimed invention, either expressly or inherently, in detail sufficient to enable one of ordinary skill in the art to practice the claimed invention without undue experimentation. I understand that inherency requires that a prior art reference necessarily functions in accordance with, or includes, the claimed limitation(s). I understand that to determine what a prior art reference discloses for the purposes of evaluating novelty, it is important and appropriate to read the reference in combination with the knowledge of those of ordinary skill in the art.

⁵⁴ 35 U.S.C. § 102.

79. I understand that patent law distinguishes novelty from obviousness and that obviousness exists if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which the claimed subject matter pertains.⁵⁵

80. I understand that in determining the obviousness of the claim(s) of a patent, one should consider:

- a. the scope and content of the prior art relied upon by the party alleging invalidity of the patent;
- b. the differences, if any, between each claim of the patent and the prior art; and
- c. the level of ordinary skill in the pertinent art at the time the invention of the patent was made; and
- d. whether the prior art enabled a person of ordinary skill in the art to make and use the invention claimed.

I understand that one must also consider such objective considerations as commercial success, long-felt but unresolved need, failure of others to solve the problem, acquiescence in the patent by others, and whether the same or similar inventions were made independently by others prior to or at about the same time as the invention claimed in the patent.

81. I understand that the test of obviousness is whether the claimed invention, as a whole, would have been obvious to one of ordinary skill in the art as of the date of the invention in light of the prior art. To establish obviousness under this test, I understand that one must show by clear and convincing evidence that a person of ordinary skill in the art at the time of invention, confronted by the same problem as the inventor and with no knowledge of the claimed invention, would select the recited

⁵⁵ 35 U.S.C. § 103.

elements from the prior art and combine them in the claimed manner. In other words, one must avoid the use of hindsight and instead identify in the art prior to the invention some suggestion or motivation, before the invention itself was made, to make the new combination.

82. I understand that the motivation to combine prior art references need not be expressly stated in the prior art, but that it may be found, for example, in the nature of the problem to be solved, or may come from the knowledge of those skilled in the art. The motivation, suggestion, or teaching to combine need not be explicit, it may be implicit from the prior art as a whole, rather than expressly stated in the reference itself. I further understand that the test for an implicit showing is what the combined teachings, knowledge of one of ordinary skill in the art, and the nature of the problem to be solved as whole would have suggested to those of ordinary skill in the art.

VIII. SUMMARY OF OPINIONS

83. Based on my review of the materials listed in Exhibit B, as well as my experience, study and training in the field of computer science and intrusion detection over the course of my career, I have formed the following opinions regarding the lack of novelty of the inventions claimed in the patents-in-suit. For convenience, I have provided a summary of the claims that are rendered invalid as anticipated or obvious by each different prior art publication or system in Exhibit JJ.

- Prior to the filing date of the original application resulting in the patents-in-suit, the subject matter described in the claims listed in Exhibit JJ was described in the printed publications listed on Exhibits C-F and V regarding NSM, DIDS, ISM, GrIDS, and *Network Intrusion Detection 1994 (NID 1994)*; it was also known or used by others in the United States with regard to the

NSM, DIDS, and GrIDS systems; it was also known by others in the United States with regard to the ISM design.

- Prior to the filing date of the original application resulting in the patents-in-suit, the subject matter described in the claims listed in Exhibit JJ was described in the printed publications listed on Exhibits G-L and U regarding JiNao, *Emerald 1997* (both alone and in conjunction with *Intrusive Activity 1994* and *NIDES 1994*), *Emerald – Conceptual Design 1997*, *Live Traffic Analysis, Network NIDES*, and the *Feather thesis*.
- Prior to the filing date of the original application resulting in the patents-in-suit, the subject matter described in the claims listed in Exhibit JJ was known or used by others in the United States with regard to the HP OpenView, NetStalker, NetRanger, and RealSecure systems; it was also described in the printed publication NetRanger User's Guide Version 1.3.1.
- Prior to the filing date of the original application resulting in the patents-in-suit, the subject matter described in the claims listed in Exhibit JJ was described in the printed publications listed in Exhibits Q, S, and T “Network Level Intrusion Detection,” “Stake Out Network Surveillance White Paper,” and “AIS.” It was also described in an application for patent filed in the United States that issued as U.S. Patent 5, 825,750 as listed in Exhibit R. Furthermore, it was known or used by others in the United States with regard to the Stake Out system. It was also known by others in the United States with regard to the “Network Level Intrusion Detection” and “AIS” systems.

84. Based on my review of the materials listed in Exhibit B, as well as my experience, study and training in the field of computer science and intrusion detection over the course of my career, I have formed the following opinions regarding the obviousness of the inventions claimed in the patents-in-suit:

- The difference, if any, between the prior art and the subject matter described in the patents-in-suit would have been obvious to one of skill in the art as of November 9, 1997. Exhibit JJ provides a list of the claims rendered obvious by each particular printed publication, patent or system, in light of the teachings of the references indicated in Exhibits C-V.
- The lack of government use and absence of a long-felt need satisfied only by the claimed inventions further demonstrates that the subject matter of the asserted claims would have been obvious to one of ordinary skill in the art as of November 9, 1997.

85. Based on my review of the materials listed in Exhibit B, as well as my experience, study and training in the field of computer science and intrusion detection over the course of my career, I have formed the opinion that the patents-in-suit are invalid because the inventors failed to disclose their best mode.

IX. STATE OF THE RELEVANT ART AS OF THE FILING DATE

A. HISTORY OF NETWORK MONITORING AND INTRUSION DETECTION

86. Network monitoring for the purpose of detecting suspicious activity has its roots in both the general-purpose computer monitoring field and the field of host-based intrusion detection. In general, computer management and security techniques are used to preserve the security and integrity of the volumes of data and applications that may reside on different computer systems. As computer networks have proliferated and the volume of data stored on computers has multiplied, the need for increased security for all these networks and data has continued to increase in importance.

87. In the 1970s, most computers systems were centralized. Typically, a mainframe processor would be connected to a group of relatively dumb terminals

allowing for data communication with the mainframe. However, by the 1980s, significant changes had occurred in computer systems and data communications. Microprocessor computers appeared, offering significant price and performance advantages over mainframes. Communication capabilities between computers also increased significantly, due to the growth of local area networks (LANs) connecting the growing numbers of microprocessor-based computers. High-speed wide area networks (WANs) also emerged to link LANs together. Networks began to proliferate. In particular, the Internet, developed from U.S. Government research that began in 1969, became increasingly important in the 1980s.

88. As networks proliferated, so did the need to manage those networks and monitor them for problems. By the late 1980s, the Internet Activities Board (IAB) recognized the need to monitor and manage the Internet, and began developing a set of standards to implement these goals. Like other Internet standards, these network management standards were documented as Request for Comments (RFCs). The first such standard, Simple Network Management Protocol (SNMP), has achieved widespread use as a tool to monitor the status of networks for suspicious activity such as network problems and faults. SNMP was designed to allow data from each individual piece of the network, such as routers and gateways responsible for directing network traffic, to communicate its status to centralized managers, creating a global view of the status of a network.

89. A wide variety of different commercial network monitoring tools and systems, such as HP OpenView, have been developed since the early 1990s to assist network administrators in the increasingly complex task of monitoring the status of networks.

90. The growth of the computer industry also led to growth in the field of computer security. Many different types of systems have been developed to try and solve

different problems related to computer security. For example, as detailed in the report of Frederick Avolio, beginning in the 1980s firewalls were developed to help block undesirable network traffic from networks. In addition, again beginning in the 1980s, so-called intrusion detection systems (“IDS”) were designed to detect, and in some cases thwart, unwanted attempts to infiltrate or access a computer or computer network. An “intrusion” can refer to any type of anomalous, illicit, or prohibited activity. An intrusion may originate from an external threat, or misuse by an internal user.

91. I am uncertain as to when the phrase “intrusion detection” was first coined, but intrusion detection as a field has grown in importance along with the growth of the more general computer security field. The first intrusion detection systems in the early 1980s naturally focused upon protecting the centralized, multi-user computer systems that existed at the time. These early intrusion detection systems typically used audit-based analysis sources (often referred to as “host-based” systems because the source of the data is internal to each host being monitored). An audit trail is a record showing who has accessed a computer system and what operations have been performed during a given period of time. An audit trail may track basic operating system functions, such as system calls and processes performed, or it may track application usage or data access.⁵⁶ I will attempt to provide a brief history of how the intrusion detection field progressed along with the growth of computer networking generally to eventually include network-centric analyses as well as host-based data sources.

92. The United States government, including the various branches of the military, has been involved since the dawn of the computer security field in funding research to protect computers and computer data.⁵⁷ The initial development of computer

⁵⁶ See S. Garfinkel and G. Spafford, PRACTICAL UNIX & INTERNET SECURITY at 289-92 (O’Reilly and Assoc. 2nd ed. 1996).

⁵⁷ Much of my own research and work, as well as much of the other research and systems addressed in this report, was partly or wholly funded by the U.S. government. The

auditing itself can in part be traced back to the U.S. government. In the 1970's, the U.S. Department of Defense ("DOD") funded an extensive research effort to provide computer system security for the processing of classified information. As part of this program, the DOD created a policy for implementing certain auditing functions to satisfy certain security goals, including "to allow the review of patterns of access" and "to allow the discovery of both insider and outside attempts to bypass protection mechanisms, and to assure that "attempts to bypass the protection will be recorded and discovered...".⁵⁸

93. James P. Anderson is often cited as the first person to propose that computer auditing mechanisms could be used by computer security personnel to track internal and external penetrations and misfeasance.⁵⁹ After Anderson's early work in 1980, several other early intrusion detection systems were developed which focused upon the analysis of audit trail information. Some of the initial analysis of audit trails was actually performed by hand, with security personnel manually reviewing audit trails to look for problems and patterns. Eventually, given the labor-intensive nature of such work, the industry began to move towards more automated systems for reviewing audit trials to detect misuse of computer systems.

94. The 1980s also saw a shift in computing generally away from more mainframe and centralized computing to networks. As networks proliferated, network monitoring tools were developed to look for network faults and communication errors. Eventually the intrusion detection community followed this trend, and intrusion detection systems began focusing upon network traffic and network sources for attack. This shift towards incorporating network-based sources of data for intrusion detection occurred

patents-in-suit also state that the inventions were made with government support provided by DARPA.

⁵⁸ See R. Bace, *INTRUSION DETECTION* at 11.

⁵⁹ See J. Anderson, *Computer Security Threat Monitoring and Surveillance*, Washington, PA, James P. Anderson Co., 1980.

much earlier than the alleged inventions in the patents-in-suit, however. By 1990, a research group at the University of California, Davis, of which I was an integral part, had already implemented an intrusion detection system using network packets as a data source.

95. The computer industry as a whole has grown enormously in the past 25 years. The field of computer security has also grown drastically in that time. While the intrusion detection community has certainly grown from its initial beginnings in the late 1980s, it has remained a rather small, tight-knit group of people. I attended numerous conferences on intrusion detection in the 1990s in particular, and there was certainly a core group of people in attendance who were well-known to each other.

96. Cooperation and collaboration among practitioners in the intrusion detection field has often occurred. For example, the Common Intrusion Detection Framework, or CIDF, was a joint effort by many of the researchers funded on a Defense Advanced Research Projects Agency (DARPA) Information Technology Office (ITO) program led by Teresa Lunt. The DARPA Program Manager, Teresa Lunt, and the DARPA researchers recognized that a lot of researchers were attempting to solve similar problems in intrusion detection, so in order to reduce duplicative work and to foster interoperability between the technology components and systems being developed by the researchers, they decided to jointly develop CIDF. For example, during the DARPA Principal Investigator's (PI) meeting in Savannah in February of 1997, the CIDF effort's proposed vision was: "To develop standards so that all ITO funded Intrusion Detection Systems can demonstrably inter-operate." The kick-off slides further state: "Six projects are, at least in part, developing an 'architecture or system for other components,'"

including Boeing, SRI, UC Davis, Purdue, GE, and North Carolina (MCNC and NCSU).⁶⁰

97. Over the next several months the CIDF members developed a series of proposals for a common network protocol, message syntax and semantics, and APIs to support interoperability between their components. The CIDF members actively worked together, communicating regularly over email (a mailing list was established), meeting at computer security related conferences, and meeting during a summer 1997 DARPA PI meeting held at SRI. Over the summer and early fall of 1997 several groups submitted proposals to the mailing list covering different aspects of CIDF. For example, the EMERALD/JiNao team submitted a proposal in July presenting a protocol for event components, analysis components, and countermeasure components to communicate with each other.⁶¹

98. In addition to Teresa Lunt's DARPA ITO program, other DOD-related efforts proclaimed they would support CIDF as well. For example, the DARPA Information Assurance program led by Sami Saydjari, of which I was part, also planned to use CIDF to support communication between their intrusion detection components.

99. Eventually the DARPA researchers attempted to broaden the effort and make CIDF an endorsed standard, so they took their proposed design to the Internet Engineering Task Force (IETF). While not endorsing the proposed CIDF solution, IETF did endorse the goal of developing interoperability standards between intrusion detection components and formed the Intrusion Detection Working Group with the goal of

⁶⁰ Common Intrusion Detection Framework "APIs" Presentation by Stuart Staniford-Chen 2/26/97 [SYM_P_0071471-SYM_P_0071481].

⁶¹ Email from P. Porras to CIDF, subject "Interface Spec: SRI, MCNC/NCSU," date 22 Jul 1997 [SYM_P_0500624-SYM_P_0500639].

developing RFC standards, the most prominent of which was called the Intrusion Detection Message Exchange Framework (IDMEF).

100. SRI was certainly well-known as an important group in the intrusion detection community. Similarly, a group of people including myself from the Computer Science Laboratory at UC Davis were also well-known as longstanding participants in the intrusion detection community. Both SRI and UC Davis were primarily research-focused. Below I have provided a short overview of the work done by these two research groups. By no means, of course, was research in intrusion detection limited to these two groups. Other important groups doing research in intrusion detection included, for example, Purdue University, Columbia University, MIT's Lincoln Laboratory, UC Santa Barbara, Lawrence Livermore National Laboratory, and Los Alamos National Laboratory.

101. The U.S. government funded a great deal of research in intrusion detection, but also actually used such systems to protect government assets, and I have provided a short history of some of these systems as well. In addition, I have provided a short history of some of the first commercial intrusion detection systems.

1. SRI's history in intrusion detection

102. SRI, in conjunction with various government research efforts, has worked and published in the IDS field for several decades. Much of this published work involves a project developed for and funded primarily by the US Navy and the US Air Force that has undergone three different evolutions over time: IDES, NIDES, and EMERALD. As the inventors themselves have explained:

Our earlier intrusion-detection efforts in developing IDES (Intrusion Detection Expert System) and later NIDES (Next-Generation Intrusion Detection Expert System) were oriented toward the surveillance of user-session and host-layer activity. This previous focus on session activity within host boundaries is understandable given that the primary input to intrusion-detection tools, audit data, is produced by mechanisms that tend to be locally administered within a

single host, or domain. However, as the importance of network security has grown, so too has the need to expand intrusion-detection technology to address network infrastructure and services. In our current research effort, EMERALD (Event Monitoring Enabling Responses to Anomalous Live Disturbances), we explore the extension of our intrusion-detection methods to the analysis of network activity.⁶²

103. Different authors at SRI, including the named inventors for the patents-in-suit, published extensively on IDES, NIDES, and EMERALD prior to filing the '338 patent.⁶³

104. In the late 1980s under a US Navy government contract, SRI began working on IDES – an audit-based system to observe computer behavior and learn to recognize “normal” behavior and deviations from expected behavior.⁶⁴ IDES used statistical “profiles” of past behavior to describe normal behavior, with the data involved being drawn from audit records. Rare behavior was referred to as “anomalous.” IDES eventually also included a rule-based system to detect known security violations.

105. Several other companies and researchers in the late 1980s also developed intrusion detection systems that used statistical methods to find anomalies in audit records. Examples include: (1) Haystack from Tracor Applied Sciences, Inc. / Haystack Laboratories, principally written by Steve Smaha, (2) Multics Intrusion Detection and Alerting System (MIDAS) from the National Computer Security Center (NCSC); (3) Network Audit Director and Intrusion Reporter (NADIR) from Los Alamos National Laboratory and (4) Wisdom and Sense from Los Alamos National Laboratory.⁶⁵

⁶² P. Porras and A. Valdes, *Live Traffic Analysis* at 3.

⁶³ See Appendix ---, listing 20 different SRI publications on IDES, NIDES, and EMERALD, all dated more than one year prior to November 9, 1998.

⁶⁴ H. Javitz and A. Valdes, *The SRI IDES Statistical Anomaly Detector*, <http://www.sdl.sri.com/papers/stats91/>, May 1991.

⁶⁵ See Rebecca Bace, “Intrusion Detection” (Macmillan Tech. Pub. 2000) (“Bace”), at 103-07.

106. Under additional contracts from the US Navy, SRI's "next-generation" of the IDES project (NIDES) extended SRI's work on statistical profiling.⁶⁶ IDES and the original NIDES were primarily host-based systems, deriving their information from audit data.⁶⁷ However, in the early 1990s researchers were increasingly focusing upon the use of network traffic in IDS, including the government agencies funding SRI's NIDES project. In 1995, SRI published D. Anderson, T. Frivold and A. Valdes, "Next-generation Intrusion Detection Expert System (NIDES) A Summary," May 1995 ("Network NIDES"), which advocated for, and explained the ease of, expanding the NIDES system to "Network NIDES." Network NIDES was designed to extend NIDES to cover network monitoring, and specifically described extending NIDES profiles to directly monitor network packets.⁶⁸ *Network NIDES* also explained that NIDES could be extended to support hierarchical monitoring via multiple monitors.⁶⁹

107. Under additional government contracts from the US Navy and the US Air Force, SRI incorporated these changes into the successor to NIDES – the EMERALD system. By December 1996 SRI had published a conceptual overview of the EMERALD

⁶⁶ R. Jagannathan et al. (including A. Valdes), *System Design Document: Next-Generation Intrusion Detection Expert System (NIDES)*, March 9, 1993 at 55.

⁶⁷ T. Lunt, *Detecting Intruders in Computer Systems*, 1993 Conference on Auditing and Computer Technology at 1-2.

⁶⁸ *Network NIDES* at 31 ("A network monitoring process could be incorporated into NIDES to read network packets and produce canonical NIDES audit records for analysis."). Note that the decision to "extend" NIDES to include network packet monitoring occurred some four years after U.C. Davis researchers had already published and implemented such an idea for their NSM IDS system. SRI and P. Porras in particular were certainly aware of the NSM project, since it is discussed and cited in *Emerald 1997*. See *Emerald 1997* at 364 and 365 [7].

⁶⁹ *Network NIDES* at 31 ("In addition, the NIDES architecture could be extended to support multiple cooperative NIDES processes that would each be responsible for a local domain, with a higher-level NIDES process responsible for the network that supports all the local domains.").

system,⁷⁰ and by October 1997, SRI had fully disclosed the EMERALD system in the *Emerald 1997* publication. *Emerald 1997* implemented the changes suggested in *Network NIDES* – specifically, monitoring network traffic via monitoring packets, and a system of hierarchical monitors.

108. On November 9, 1998, P. Porras and A. Valdes filed the patent application which matured into the '338 patent. The '203, '212, and '615 patents all stem from the original '338 patent disclosure. These patents relate to the EMERALD system.

109. After the '338 patent filing, SRI continued to work on the EMERALD system. SRI also continued to receive additional government contracts for work on EMERALD.⁷¹ As discussed elsewhere in my report, SRI also filed additional patent applications on EMERALD-related work not disclosed in the patents-in-suit.

2. UC Davis's history in intrusion detection

110. Various researchers at the University of California, Davis have also contributed substantially to the IDS field over several decades. Beginning in the late 1980s, UC Davis developed the Network Security Monitor (NSM) for Lawrence Livermore National Laboratory (LLNL), which is widely credited as being the first IDS that monitored network traffic.⁷² LLNL secured funding from the Department of Energy (DOE) and approached UC Davis's Professor Levitt to be the Principal Investigator (PI) on the project. Professor Levitt was well known for his work in computer security, was

⁷⁰ P. Porras and P. Neumann, *EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances, Conceptual Overview*, <http://www.sdl.sri.com/papers/emerald-position1/> (December 18, 1996).

⁷¹ See SRI 005524-25, discussing contract F30602-99-C-0149, which SRI has identified as being associated with Emerald. See SRI's Second Supplemental Response to Symantec's Interrogatories Nos. 5 and 8 (#8).

⁷² R. Bace, INTRUSION DETECTION at 19; SRI 097123-86 at 097167-68, 097178.

previously Director of the SRI Computer Science Laboratory. While Professor Karl Levitt was the PI, I was the primary developer for the NSM project.

111. The UC Davis work on NSM was carried out from 1988 through about 1995, although the primary work was performed from 1988 to 1993. Following 1993 the Air Force took a copy of the NSM and developed it further under the name Automated Security Incident Measurement (ASIM), and along with the Common Intrusion Detection Director (CIDD), continues to develop it to this day. Also around 1993 LLNL took their copy of NSM and further developed it under the name Network Intrusion Detector (NID) for the DOE and provided NID to the Defense Information Systems Agency (DISA) under the name Joint Intrusion Detection System (JIDS).

112. The first major system to integrate host and network-based monitoring was the Distributed Intrusion Detection System (DIDS), introduced in the early 1990s. In addition to integrating network and host information, DIDS also implemented a monitoring hierarchy, and correlated evidence of activity across a network. The first phase of DIDS from 1990-1992 was funded by the United States Air Force, with Lawrence Livermore National Laboratory (LLNL) as the prime contractor and UC Davis and Haystack Laboratories as subcontractors.

113. In 1993 DIDS was handed over to Trident Data Systems (TDS) to develop into a supported product to be deployed across the Air Force. TDS continued to work on DIDS until about 1996, at which point TDS and the Air Force redirected their focus towards Automated Security Incident Measurement (ASIM), along with the Common Intrusion Detection Director (CIDD).

114. UC Davis also conceived and developed the Graph-based Intrusion Detection System (GrIDS). UC Davis began development of GrIDS in 1995 as part of an initial 1993 DARPA contract to develop intrusion detection systems for very large networks. UC Davis continued the GrIDS work under a 1996 contract, and continued to

work on it until about 1999. The GrIDS system was running publicly at UC Davis by approximately early 1997.

3. Government systems' history in intrusion detection

115. The United States government was one of the earliest sponsors and adopters of intrusion detection technologies and systems. For example, the MIDAS intrusion detection system was running on the NSA's Dockmaster computer by the late 1980s.⁷³ However, wide-spread adoption of intrusion detection systems within the government did not occur until the deployment of the NSM (later called the ASIM sensor) throughout the Air Force in the early to mid-1990s. In addition to the Air Force's ASIM effort, the NSM was used by the Defense Information Systems Agency (DISA), initially as just NSM and then under the name Joint Intrusion Detection System (JIDS), as well as by the Department of Energy in general and Lawrence Livermore National Laboratory in particular under the name Network Intruder Detector (NID).

116. With the availability of commercial network-based intrusion detection systems in the mid to late 1990s, the government started adopting commercial IDSs as well. In particular, NetRanger was extensively evaluated by the DOD and the Air Force bought a number of systems.⁷⁴ ISS's RealSecure also became very popular in the government, and I recall one of DARPA's Grand Challenge efforts centering around the analysis of sensor logs from RealSecure. When meeting with government officials in the early 2000s, however, I was finding more and more government sites adopting the open source system called Snort. During an Intelligence Community (IC) Principal Investigators meeting in 2005, the only intrusion detection system I heard people using (besides some of their own research efforts) was Snort.

⁷³ B. Mukherjee et al., "Network Intrusion Detection," IEEE Network, May/June 1994.

⁷⁴ See Expert Report of Daniel Teal.

117. It is fair to say that today there is a wide mix of government off the shelf (GOTS), commercial off the shelf (COTS), and open source intrusion detection systems being used within the government, and most of these systems are network-based intrusion detection systems.

4. Commercial systems' history in intrusion detection

118. Haystack Laboratories was one of the first commercial companies developing intrusion detection systems. While beginning their efforts with government contracts (the Haystack intrusion detection system for the Air Force in the late 1980s and DIDS in the early 1990s), they eventually developed their own commercial intrusion detection systems around 1993 with Stalker (a host-base IDS). Stalker was quickly followed by NetStalker (a network-based IDS) and WebStalker (an application-based IDS). In the mid-1990s ISS, which started with a vulnerability scanner, developed a network-based intrusion detection system called RealSecure, and WheelGroup, formed by several people who had worked with the NSM in the Air Force, developed NetStalker. During the late 1990s, in part fueled by the Internet stock bubble, a large number of commercial intrusion detection companies formed. Today, intrusion detection systems are part of the product lines of many companies including Cisco, Juniper, and Symantec.

119. See also the expert report of Frederick Avolio for a history of related network entities such as packet filters and firewalls, and the expert report of Daniel Teal for the history of NetRanger.

X. LACK OF NOVELTY

120. The following sections discuss the features and functionality of each of the main prior art publications and systems that in my opinion demonstrate that some or all of the asserted claims are invalid as anticipated.

A. PRIOR ART PUBLICATIONS AND SYSTEMS

1. NSM

121. The Network Security Monitor (NSM) was a network intrusion detection system first developed in the late 1980s at UC Davis that proliferated to many different organizations and has undergone many different permutations since its inception. Many different publications have described various aspects of the NSM system, as listed in Exhibit C. This report focuses on the NSM as it was described in: L.T. Heberlein, G.V. Dias, K. N. Levitt, B. Mukherjee, J. Wood, D. Wolber, *A Network Security Monitor*, Proc. 1990 Symposium on Research in Security and Privacy, pp. 296-304, May 1990 (“NSM 1990”).

122. The NSM was designed to protect a network from attack. The target system to protect included “a number of computers (including devices such as file servers, name servers, printers, etc.) and a LAN through which the hosts are interconnected.”⁷⁵ The LAN included “the wire, bridges, routers, and gateways.”⁷⁶ Thus, NSM received packets from a network entity such as a router or a gateway.

123. The NSM project was inspired, to a large extent, by the attacker that Cliff Stoll tracked across the globe – that attacker also crossed DOE networks. Few if any of the penetrated computers generated audit trails that could be processed by intrusion detection systems, but in each case the attacks did generate network packets. The NSM was therefore developed to process packets as the data source. The NSM processed TCP/IP packets on an Ethernet, but *NSM 1990* clearly states that the approach should work on other protocols and network topologies.

⁷⁵ *NSM 1990* at 297.

⁷⁶ *NSM 1990* at 297.

124. While the NSM used network traffic for its data source, *NSM 1990* clearly states that it uses statistical detection analysis concepts from host-based intrusion detection systems such as SRI's IDES.⁷⁷ This was not unusual – statistical analysis techniques, including SRI's statistical profiling techniques, were widely known and had been widely shared with the intrusion detection community. Specifically, the NSM used a combination of anomaly-based statistical detection using profiles and signature-based detection using rules to detect potentially intrusive behavior.

125. The NSM had the following salient features:

Designed to protect hosts (including file servers, name servers, printers, etc.) and a LAN (defined as the wire, bridges, routers, and gateways).

Implemented for the TCP/IP protocol on a CSMA/CD network.

Extendable to other network protocols and network topologies.

Processed network packets.

Modeled several different types of objects including hosts, data paths, and connections.

Monitored network connections and network data volume.

Used both anomaly detection and signature detection to determine if an object was behaving suspiciously.

Responded to detected suspicious activity by reporting objects to a security officer and potentially modifying analysis.

Stated plan to extend the approach to a distributed environment with multiple network monitors exchanging information.

126. The NSM modeled several subjects that shared a hierarchical relationship with each other: source objects, source-destination objects, source-destination-service objects, and connections. These subjects, examples of each, and their relationships are shown in Figure 1 of Exhibit W. The following discussion refers to the number labels in Fig. 1.

⁷⁷ *NSM 1990* at 296.

127. The top-level objects are source objects (Src) – hosts that initiate connections. Figure 1 shows two source objects: Venus (1) and Mars, and provides examples of the types of measurements that would be propagated through the NSM. The NSM tracked two measurements for each source object: the number of packets (Pkts) for connections initiated by the host and the number of hosts to which this host tried to connect.

128. The second level objects are source-destination objects (Src-Dst) – these represent pairs of communicating hosts. For example, (2) represents activity in which the host Venus initiated connections to the host Mercury. The NSM tracked two measurements for each source-destination object: the number of packets over connections originating from the source to this destination and the number of different services over which the source communicated to the destination.

129. The third level objects are source-destination-service objects (Src-Dst-Srv) – pairs of hosts communicating over a specific service such as Telnet, FTP, or Sendmail. For example, (3) represents activity from connections initiated by the host Venus to the host Mercury over the Telnet service. The NSM tracked two measurements for each source-destination-service object: the number of packets over connections originating from the source to this destination over this service and the number of different connections between the pair of hosts and this service.

130. The fourth level objects are the connections themselves (Connections). For example, (4) represents a single connection initiated by Venus to Mercury over the Telnet service. The NSM tracked two measurements for each connections object: the number of packets over the connection and the number of bytes in those packets.⁷⁸

⁷⁸ NSM 1990 mentions only the number of packets, but the NSM system itself tracked both the number of packets and bytes.

131. To determine if a subject was behaving maliciously and a security officer should be alerted, the NSM used three different methods: anomaly detection on the metrics for each subject, signatures on the metrics for each subject, and anomaly detection on the existence of data paths.

132. NSM's anomaly detection used the values from past instances of objects to build a profile of that object's behavior, compare new instances of the object against the profile, and report objects that were behaving anomalously, as shown in Figure 2. NSM's anomaly detection process began by collecting measurements of past historical activity (see top level of Figure 2). For example (1) shows a past instance of a subject that had 121 packets, and (2) shows a past instance of a subject that had 71 packets.

133. NSM took these past instances and built a probability distribution that measured how frequently past values occurred – this was the *profile* for the measurement. (see second level of Figure 2). For example, in Figure 2 the X-axis for the probability distribution represents the number of packets seen for a measurement – on the far left the first bar represents instances which had 0-9 packets, the next one represents instances that had 10-19 packets, and so on (3). The Y-axis represents how frequently that measurement was observed (4). For example, the bar on the far left of the X-axis (0-9 packets) is very small, indicating that very few past measurements had this value (roughly 1%), while the bar for objects with 50-59 packets were the most common, occurring almost 20% of the time. This probability distribution was exponentially aged.

134. The NSM then compared this long-term statistical profile to recent activity. To determine if the difference between the historical and recent activity was anomalous enough to be reported, the NSM security officer set a threshold (5) for the measurement, and any future measurement that fell in a range (or bin) with a frequency below that threshold was reported. This is functionally identical to comparing the

difference between the profiles to the “score threshold” of the patents-in-suit.⁷⁹ To implement this approach, the NSM used the concept of a mask (6).

135. A mask (6) is an object that by default, when presented with a value, blocks that value. However, holes (7) can be put into the mask that will allow some values to flow through the mask. Values that flow through the mask are reported. The NSM created an anomaly mask (6) (see third level of Figure 2) and cut holes (7) in the mask corresponding to each bin in the profile that fell below the threshold (5). In the Fig. 2 example, there are holes in the mask for bins representing the ranges 0-9, 10-19, and 80-89. This anomaly mask became the mechanism used to determine if a future measurement was anomalous enough to be reported (see fourth level of Figure 2). For example, (8) and (10) represent measurements of future instances that are tested against the anomaly mask. Instance (8) has 90 packets, and when tested against the mask, the mask blocks this (9) from being reported. However, Instance (10) has 83 packets, and when tested against the mask falls through the hole (11) to become an anomalous event (12) that is reported to the NSM security officer (16).

136. NSM’s measurements of recent activity were statistical descriptions of recent activity. I understand that SRI has claimed that a current snapshot of recent activity is not a short-term statistical profile. That is incorrect. In fact, Mr. Valdes, one of the named inventors, testified that a short-term profile could be a single measurement or event.⁸⁰

137. In addition, under Symantec’s alternative claim construction, the NSM would satisfy the requirement that the long-term statistical profile was an exponentially aged probability distribution of historical values. Although NSM’s short-term statistical description was not aged, *NSM 1990* specifically pointed the reader to SRI’s IDES

⁷⁹ ‘338 col. 6:59-7:3.

⁸⁰ Valdes Tr. 376-377.

project, which became the NIDES project. The NIDES project had extensive publications about the algorithms for such a short-term statistical profile, and one of skill would have been motivated to combine the two systems.⁸¹

138. In addition to detecting anomalous activity, NSM applied a set of rules, or signatures, to look for specific patterns of behavior. *NSM 1990* states that signatures are particularly important when the intrusion detection system is first deployed and has not had time to build up useful profiles. To implement the signatures, the NSM used the same masking technology developed for the anomaly detection, as shown in Fig. 2. Examples of rules mentioned in the paper include:

- Looking for connections with small numbers of packets indicating a possible failed login.
- Looking for a host connecting to large numbers of other hosts indicating a host probing the network.

139. The final analysis method mentioned in *NSM 1990* is another form of anomaly detection, but instead of measuring the behavior of an object (e.g., by measuring the number of packets associated with the object), the NSM measured the probability of the object just existing. This analysis was only applied to connections, and the measurement is the probability of seeing a connection from a host **A** to host **B** via service **C**. This triple (which is essentially the Source-Destination-Service class of objects) is called a “data path.” When a new connection was observed, the NSM looked up a profile to see if this data path had been used in the past. If not, then the connection was considered anomalous. In *NSM 1990* the data path profile was relatively simple: a data path was considered normal if it had been used at least once in the previous two weeks. Subsequent implementations used an exponentially aged distribution of previously-observed connection patterns.

⁸¹ See, e.g., *Statistical Methods*.